arXiv:math/0107033v1 [math.PR] 4 Jul 2001

Information flow on trees

Elchanan Mossel
INRIA Rocquencourt and Microsoft Research

Yuval Peres
Hebrew University and U.C. Berkeley

February 1, 2008

Abstract

Consider a tree network T, where each edge acts as an independent copy of a given channel M, and information is propagated from the root. For which T and M does the configuration obtained at level n of T typically contain significant information on the root variable? This problem arose independently in biology, information theory and statistical physics.

- For all b, we construct a channel for which the variable at the root of the b-ary tree is independent of the configuration at level 2 of that tree, yet for sufficiently large $B > b$, the mutual information between the configuration at level n of the B-ary tree and the root variable is bounded away from zero. This is related to certain secret-sharing protocols.
- We improve the upper bounds on information flow for asymmetric binary channels (which correspond to the Ising model with an external field) and for symmetric q-ary channels (which correspond to Potts models).
- Let $\lambda_2(M)$ denote the second largest eigenvalue of M, in absolute value. A CLT of Kesten and Stigum (1966) implies that if $b|\lambda_2(M)|^2 > 1$, then the *census* of the variables at any level of the b-ary tree, contains significant information on the root variable. We establish a converse: if $b|\lambda_2(M)|^2 < 1$, then the census of the variables at level n of the b-ary tree is asymptotically independent of the root variable. This contrasts with examples where $b|\lambda_2(M)|^2 < 1$, yet the *configuration* at level n is not asymptotically independent of the root variable.

1 Introduction

Consider a process in which information flows from the root of a tree T to other nodes of T. Each edge of the tree acts as a channel on a finite alphabet $\mathcal{A} = \{1, \dots, k\}$. Denote by $\mathbf{M}_{i,j}$ the transition probability from i to j, and by M the random function (or channel) which satisfies for all i and j that $\mathbf{P}[M(i) = j] = \mathbf{M}_{i,j}$. Let $\lambda_2(M)$ denote the eigenvalue of $\mathbf{M}$ which has the second largest absolute value ($\lambda_2(M)$ may be negative or non-real). At the root ρ one of the symbols of $\mathcal{A}$ is chosen according to some initial distribution. We denote this (random) symbol by σ_ρ. This symbol is then propagated in the tree as follows. For each vertex v having as a parent v', we let $\sigma_v = M_{v',v}(\sigma_{v'})$, where the $\{M_{v',v}\}$ are independent copies of M. Equivalently, for a vertex v, let

v' be the parent of v, and let $\Gamma(v)$ be the set of all vertices which are connected to ρ through paths which do not contain v. Then the process satisfies:

$$\mathbf{P}[\sigma_v = j | (\sigma_w)_{w \in \Gamma(v)}] = \mathbf{P}[\sigma_v = j | \sigma_{v'}] = \mathbf{M}_{\sigma_{v'}, j}.$$

It is very natural to study this process in the context of biology, statistical physics and communication theory. See [10],[22] and the references there for more background.

Let $d(,)$ denote the graph-metric distance on T, and $L_n = \{v \in V : d(\rho, v) = n\}$ be the n'th level of the tree. For $v \in V$ and $e = (v, w) \in E$ we denote $|v| = d(\rho, v)$ and $|e| = \max\{|v|, |w|\}$. We denote by $\sigma_n = (\sigma(v))_{v \in L_n}$ the symbols at the n'th level of the tree. We let $c_n = (c_n(1), \ldots, c_n(k))$ where

$$c_n(i) = \#\{v \in L_n : \sigma(v) = i\}.$$

In other words, c_n is the *census* of the n'th level. Note that both $(\sigma_n)_{n=1}^{\infty}$ and $(c_n)_{n=1}^{\infty}$ are markov chains.

For distributions P and Q on the same space the total variation distance between P and Q is

$$D_V(P, Q) = \frac{1}{2} \sum_{\sigma} |P(\sigma) - Q(\sigma)|. \tag{1}$$

Definition 1.1. *The reconstruction problem for T and M is* **solvable** *if there exist $i, j \in \mathcal{A}$ for which*

$$\lim_{n \to \infty} D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) > 0, \tag{2}$$

where $\mathbf{P}_n^{\ell}$ denotes the conditional distribution of σ_n given that $\sigma_\rho = \ell$.

Definition 1.2. *The reconstruction problem for T and M is* **census-solvable** *if there exist $i, j \in \mathcal{A}$ for which*

$$\lim_{n \to \infty} D_V(\mathbf{P}_n^{(c),i}, \mathbf{P}_n^{(c),j}) > 0, \tag{3}$$

where $\mathbf{P}_n^{(c),\ell}$ denotes the conditional distribution of c_n given that $\sigma_\rho = \ell$.

Assume that $\mathbf{P}[\sigma_\rho = i] > 0$ for every $i \in \mathcal{A}$. Then the following conditions are equivalent, see e.g. [22]:

- The reconstruction problem for T and M is not solvable.
- The configurations σ_n are asymptotically independent of σ_ρ, i.e., $\lim_{n \to \infty} I(\sigma_\rho, \sigma_n) = 0$, where I is the mutual information operator.
- The sequence $(\sigma_n)_{n=1}^{\infty}$ has a trivial tail σ-field.
- The distribution of $\{\sigma_v\}_{v \in T}$ is an extremal Gibbs measure.

Similar conditions apply to census-solvability, where $\{c_n\}$ replace $\{\sigma_n\}$. Note that if the reconstruction problem is census-solvable it is also solvable.

The reconstruction problem was first studied in statistical physics for the Ising model on the tree, in the equivalent form as a question regarding the extremality of the free Gibbs measure for this model. See [25],[13], [5],[14], [15] and [16]. Reference [10] contains extensive background on this problem for the Ising model. In these papers the reconstruction problem for the Ising model on trees was solved. Writing $\lambda_2(M)$ for the second eigenvalue of $\left(\begin{array}{cc} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{array}\right)$ (that is, $\lambda_2(M) = 1-2\epsilon$), the reconstruction problem is solvable for the b-ary tree $T_b = (V_b, E_b)$ if and only if $b\lambda_2^2(M) > 1$. An analogous threshold for general trees was established in [10].

A few combinatorial questions arise naturally in the context of these tree processes (see also [2], [3] and [4]). In Section 2 we discuss the following question. For a channel M we ask does there exist any b such that the reconstruction problem is solvable for the infinite b-ary tree T_b and the channel M? In Theorem 2.1 we present a criteria for deciding this problem for a channel M. Using this criteria we construct a channel with the following properties:

Example 1.1. *Write M^h for the h'th iteration of the channel M. Then for all $h > 1$ there exists a channel M such that*

- *$M^h(j)$ has the same distribution for all $j \in A$.*
- *For all $h' < h$ there exist i and j such that $M^{h'}(i)$ and $M^{h'}(j)$ have different distributions.*
- *When b is large the reconstruction problem is solvable for the tree T_b and the channel M.*

This is a generalization of an example which appeared in [21] and [22]. In Section 3 we give more delicate constructions which are related to secret-sharing protocols (see [24]).

Theorem 1.2. *Let $b > 1$ be an integer and T be the 2-level b-ary tree. There exist a channel M such that for any initial distribution, σ_r and σ_∂ are independent (where σ_∂ is the configuration at the leaves of the 2-level b-ary tree), yet when B is sufficiently large, the reconstruction problem for the channel M and the B-ary tree T_B is solvable.*

It is tempting to try to find thresholds for the reconstruction problem which depend only on b and $\lambda_2(M)$. For binary symmetric channels the threshold for reconstruction is $b\lambda_2^2(M) = 1$ (this is also the threshold for census-reconstruction for general channels: see theorems 1.4 and 1.5 below).

In [22] it is shown that for some natural generalizations of the binary symmetric channel, the threshold $b\lambda_2^2(M) = 1$ is not the threshold for reconstruction. In particular, it is shown that for the Potts model (the q-ary symmetric channel), where the transition matrix is:

$$\mathbf{M} = \left(\begin{array}{llll} 1-(q-1)\delta & \delta & \dots & \delta \\ \delta & 1-(q-1)\delta & \delta & \dots \\ \vdots & \dots & \ddots & \vdots \\ \delta & \dots & \delta & 1-(q-1)\delta \end{array}\right), \qquad (4)$$

if $b\lambda_2(M) = b(1-q\delta) > 1$ and q is sufficiently large, then the reconstruction problem is solvable. Similar results hold for asymmetric binary channels. On the other hand, it is well known, see e.g.

[22], that for both families the reconstruction problem is unsolvable if $0 \le b\lambda_2(M) \le 1$. In Section 4 we improve the bound $b\lambda_2(M) \le 1$ by proving:

Proposition 1.3. *Let* $\mathbf{M}$ *be the matrix (4). Then the reconstruction problem for* T_b *and* M *is unsolvable when*

$$b\frac{(1-q\delta)^2}{1-(q-2)\delta} \le 1. \tag{5}$$

Similar results hold for asymmetric binary channels. The proofs are based on a reduction to symmetric binary channels on general trees which were analyzed in [10].

In sections 5, 6 and 7 we consider census reconstruction and show that the threshold for census reconstruction is $b|\lambda_2(M)|^2 = 1$. From [17] it follows that if $b|\lambda_2(M)|^2 > 1$, then the reconstruction problem is census solvable. We extend this result to general trees by proving:

Theorem 1.4. *Let* T *be an infinite tree and write* $\mathrm{br}(T)$ *for the branching number of the tree. Let* M *be a channel such that* $\mathrm{br}(T)|\lambda_2(M)|^2 > 1$*, then the reconstruction problem is solvable for* T *and* M*.*

The results of [17] also plays a crucial role in proving:

Theorem 1.5. *The reconstruction problem for the* b*-ary tree and the channel* M *is not census-solvable if* $b|\lambda_2(M)|^2 < 1$.

We conjecture that the result of Theorem 1.5 should hold also when $b|\lambda_2(M)|^2 = 1$, but we have verified this only for the following channels:

Theorem 1.6. *Let* M *be the* q *state Potts model, or the asymmetric Ising model, and suppose that* $b\lambda_2^2(M) \le 1$*, then the reconstruction problem is unsolvable for the* b*-ary tree and the channel* M*.*

In Section 8 we discuss some open problems.

Some of the results in this paper concern general trees. For an infinite tree T many of its probabilistic properties are determined by the branching number $\mathrm{br}(T)$. This is the supremum of the real numbers $\lambda \ge 1$, such that T admits a positive flow from the root to infinity, where on every edge e of T, the flow is bounded by $\lambda^{-|e|}$. Here $|e|$ denotes the number of edges, (including e) on the path from e to the root; $\mathrm{br}(T)^{-1}$ is the critical probability for Bernoulli percolation on T. See [18] and [10] for equivalent definitions of $\mathrm{br}(T)$ in terms of percolation, cutset sums and electrical conductance. We note that for the regular tree T_b we have $\mathrm{br}(T_b) = b$.

2 Distinguishing states by tree networks

Let M be a channel on an alphabet $\mathcal{A}$ of size k with $\lambda_2(M) = 0$. Looking at the Jordan form of $\mathbf{M}$ we see that $\mathrm{rank}(\mathbf{M}^k) = 1$, and therefore $M^k(i)$ has the same distribution for all $i \in \mathcal{A}$. Moreover, by Theorem 1.5 it follows that for all b, the reconstruction problem is not *census* solvable for the channel M and the tree T_b. Does this mean that the reconstruction problem is unsolvable for T_b and the channel M? In this section we answer this question.

Definition 2.1. *Let M be a channel on $\mathcal{A}$ and consider the minimal equivalence relation $\sim$ on $\mathcal{A}$ which satisfies: If i and j satisfy for all ℓ,*

$$\sum_{\ell' \sim \ell} \mathbf{M}_{i,\ell'} = \sum_{\ell' \sim \ell} \mathbf{M}_{j,\ell'}, \tag{6}$$

then $i \sim j$. If $i \sim j$, we say that i and j are indistinguishable*. Otherwise, we say that i and j are* distinguishable*.*

The motivation for this definition is that if $\mathbf{M}_{i,\ell} = \mathbf{M}_{j,\ell}$ for all ℓ , then after one application of the random function M, one cannot distinguish between the hypotheses that i was the input and the hypotheses that j was the input.

Theorem 2.1. *If the states i and j are indistinguishable, then there exists N such that for all $n \geq N$ and for all infinite rooted trees T:*

$$D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) = 0. \tag{7}$$

On the other hand, for every channel M, there exits b, such that for the tree T_b:

$$\inf_{n \geq 1, i \nsim j} D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) > 0. \tag{8}$$

Example 2.2. *Let $\{Z_i\}_{i=1}^{\infty}$ be an i.i.d sequence of variables such that $\mathbf{P}[Z_i = 0] = 1 - P[Z_i = 1] = p$ where $0 < p < 1$. Let $h \geq 1$ and consider the channel M defined by the markov chain $Y_i = (Z_i, \ldots, Z_{i+h})$. Thus M has state space $\{0,1\}^{h+1}$ with the product $(p, 1-p)$ probability measure. It is easily seen that all the states of M are indistinguishable. Therefore by Theorem 2.1 there exists N such that for all trees and all $n \geq N$, it holds that $D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) = 0$ (one can take $N = h$).*

Example 1.1. *Let $\{Z_i\}_{i=1}^{\infty}$ be as in Example 2.2. Set*

$$Y_i = \max_{0 \leq j \leq h} \{Z_i = \cdots = Z_{i+j} = 1\}.$$

It is easily seen that Y_i defines a channel M on the space $\{0, \ldots, h\}$. Moreover, it is clear that for all ℓ the variables $\{Z_i\}_{i \geq \ell+h+1}$ and $\{Y_i\}_{i \leq \ell}$ are independent. Therefore $\{Y_i\}_{i \geq \ell+h+1}$ and $\{Y_i\}_{i \leq \ell}$ are independent. It follows that the variables $M^{h+1}(j)$ have the same distribution for all j. Thus $\mathrm{rank}(\mathbf{M}^{h+1}) = 1$*, and $\lambda_2(M) = 0$. Writing $\mathbf{M}$:*

$$\mathbf{M} = \begin{pmatrix} p & p(1-p) & p(1-p)^2 & \ldots & (1-p)^h \\ 1 & 0 & \ldots & \ldots & 0 \\ 0 & 1 & 0 & \ldots & \vdots \\ \ldots & \ddots & \ddots & \ddots & \vdots \\ 0 & \ldots & 0 & p & 1-p \end{pmatrix},$$

we see that all the states of M are distinguishable. Therefore by Theorem 2.1, when b is sufficiently large we have for the tree T_b that $\inf_{n \geq 1, i, j} D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) > 0$. Note that the channel above is obtained by lumping the channel of example 2.2. This is a generalization of a channel appearing in [21]; see also [20].

Remarks:

- Theorem 2.1 implies that if for a channel M all the states are indistinguishable, then $\lambda_2(M) = 0$.
- Suppose that the channel M is reversible and has $\lambda_2(M) = 0$. Then looking at the diagonal form of $\mathbf{M}$, it follows that rank($\mathbf{M}$) = 1. In particular, all the states of M are indistinguishable.

Proof of Theorem 2.1: The first claim is easy. We define a new equivalence relation between states in $\mathcal{A}$. We let i and j be equivalent if there exists an N such that for all trees and all $n \geq N$ we have $D_V(\mathbf{P}^i_n, \mathbf{P}^j_n) = 0$. It is clear that this is indeed an equivalence relation and that it satisfies (6). The proof of the first claim follows.

For the proof of the second claim let $\mathcal{A}'$ be the set of equivalence classes for $\sim$. Let M' be the channel on $\mathcal{A}'$ defined as follows. For $i', j' \in \mathcal{A}'$ we choose $i \in i'$ and define

$$\mathbf{P}[M'(i') = j'] = \sum_{j \in j'} \mathbf{P}[M(i) = j]. \tag{9}$$

It is clear that (9) does not depend on the choice of i. Moreover, it is clear that for any tree T, we may couple the M tree process and the M' tree process in such a way that for all v in T we have that $\sigma_v \in \sigma'_v$. In particular, it follows that for all trees and all $i \in i'$ and $j \in j'$:

$$D_V(\mathbf{P}^{i'}_n, \mathbf{P}^{j'}_n) \leq D_V(\mathbf{P}^i_n, \mathbf{P}^j_n).$$

Therefore, in order to prove the second claim it suffices to show that there exists a b such that for the tree T_b and the channel M' we have,

$$\inf_{i,j \in \mathcal{A}', n} D_V(\mathbf{P}^i_n, \mathbf{P}^j_n) > 0. \tag{10}$$

The crucial property of $\mathbf{M}'$ we will exploit is that for all $i \neq j \in \mathcal{A}'$ we have $(\mathbf{M}'_{i,\ell}) \neq (\mathbf{M}'_{j,\ell})$. Thus in order to simplify the presentation of the proof, we will assume that the channel $\mathbf{M}$ on the alphabet $\mathcal{A}$ satisfies $(\mathbf{M}_{i,\ell})_{\ell=1}^{|\mathcal{A}|} \neq (\mathbf{M}_{j,\ell})_{\ell=1}^{|\mathcal{A}|}$ for all $i \neq j \in \mathcal{A}$

We are going to show that for every $\epsilon > 0$, there exits $b = b(\epsilon)$, and a recursive algorithm, which given the symbols at the n'th level of T_b, reconstruct the symbol at the root with probability at least $1 - \epsilon$ (uniformly for all initial distributions of the root). This implies (see [22]) that (10) holds.

We let $\delta = \min_{i,j} |(\mathbf{M}_{i,\ell}) - (\mathbf{M}_{j,\ell})|_\infty$, and assume that $\epsilon < \delta/8$. By standard results in the theory of large deviations (see e.g. [9]), there exits a positive constant C such that for all i, if X_i is the empirical distribution vector of b independent trials with distribution $\mathbf{M}_{i,\ell}$, then,

$$\mathbf{P}[\,|\,X_i - b\mathbf{M}_{i,\ell}\,|_\infty\,| > b\delta/8] \leq \exp(-Cb\delta). \tag{11}$$

Similarly, if X is the sum of b i.i.d. $\{0,1\}$ variables, each of which has the value 1 with probability ϵ and the value 0 with probability $1 - \epsilon$, then,

$$\mathbf{P}[|X - \epsilon b| > b\delta/8] \leq \exp(-Cb\delta). \tag{12}$$

We now choose b such that $\exp(-Cb\delta) \leq \epsilon/2$, and apply the following recursive algorithm in order to reconstruct the symbol at the root of the tree:

- For each vertex v, construct the empirical distribution vector X_v of the reconstructed values for the children of that vertex.
- Reconstruct at v the value i, where i minimizes the distance $|\, X_v - b\mathbf{M}_{i,\ell} \,|_\infty$.

It now follows by induction, that the probability of correct reconstruction is at least $1 - \epsilon$. Indeed if this is true for all the children of v, then by (12), with probability at least $1 - \epsilon/2$, for at least $(1 - \delta/4)b$ of the children we reconstructed the correct value. Thus by (11), with probability at least $1-\epsilon$, the vector X_v satisfies $|\, X_v - b\mathbf{M}_{i,\ell} \,|_\infty < b\delta/2$ where i is the symbol at v. It now follows that we reconstruct the correct value with probability at least $1 - \epsilon$. ■

3 The distinguishing power of tree networks

Let M be a channel on $\mathcal{A}$ and let T be a tree which consists of 2 nodes, the root r and an additional vertex v (the 1-ary 1-level tree). Suppose that $I(\sigma_r, \sigma_{\partial T}) = 0$, where ∂T are the vertices in the boundary of the tree T (in this case v), and I is the mutual information operator. The assumption $I(\sigma_r, \sigma_{\partial T}) = 0$ is equivalent to $\text{rank}(\mathbf{M}) = 1$. Therefore, all the states of M are indistinguishable and the reconstruction problem for T_b and the channel M is unsolvable for all b.

On the other hand, let T be the 2-level 1-ary tree. Example 1.1 is a channel M such that for the tree T we have $I(\sigma_r, \sigma_{\partial T}) = 0$, yet the reconstruction problem for T_b is solvable when b is sufficiently large.

It is natural to ask similar question for other finite trees T: Suppose that the channel M satisfies for any initial distribution $I(\sigma_r, \sigma_{\partial T}) = 0$, does this imply that the reconstruction problem for T_b is unsolvable for all b?

Arguing as before, it is clear that this is the case if there exists a vertex in ∂T which is at distance 1 from the root. We show that these are the only trees for which the 0-information condition implies non-reconstruction for all b.

In the following theorem we give a construction related to secret-sharing protocols [24].

Theorem 1.2 *Let $b > 1$ be an integer and T be the 2-level b-ary tree. There exists a channel M_b such that for any initial distribution $I(\sigma_r, \sigma_{\partial T}) = 0$, yet when B is sufficiently large, the reconstruction problem for the channel M and T_B is solvable.*

Definition 3.1. *Let $\mathcal{F}$ be a finite field with $q > b + 2$ elements. Let $x_1, \ldots, x_{b+1}$ be a fixed set of non-zero element of $\mathcal{F}$. We define a channel on the state space*

$$\mathcal{F}^b[x] = \{f(x) : f(x) \in \mathcal{F}[x], \deg f \le b\}.$$

Given f, take I to be a uniform variable in the set $\{1, \ldots, b+1\}$, then take $M(f)$ to be $g \in \mathcal{F}^b[x]$ which satisfy $g(0) = f(x_I)$ with probability

$$\Big|\{g \in \mathcal{F}^b[x] : g(0) = f(x_I)\}\Big|^{-1} = q^{-b}.$$

Proposition 3.1. *Let* $\{x_1, \dots, x_{b+1}\}$ *be a set of elements of* $\mathcal{F}$*. There is a one to one mapping from* $\mathcal{F}^b[x]$ *to* $\mathcal{F}^{b+1}$ *defined by*

$$f \to (f(x_1), \dots, f(x_{b+1})) \,. \tag{13}$$

The inverse map is defined by the interpolation polynomial:

$$(y_1, \dots, y_{b+1}) \to f(x) = \sum_{i=1}^{b+1} \frac{\prod_{j \neq i}(x - x_j)}{\prod_{j \neq i}(x_i - x_j)} y_i. \tag{14}$$

Lemma 3.2. *Let* T *be the 2-level* b*-ary tree. Then for any initial distribution,* $I(\sigma_r, \sigma_\partial) = 0$*.*

Proof. We show that for all $f \in \mathcal{F}^b[x]$ and $h = (h_{i,j})_{i,j=1}^b \in (\mathcal{F}^b[x])^{b^2}$ it holds that

$$\mathbf{P}[\sigma_{\partial T} = h | \sigma_r = f]$$

is independent of f. This would imply the claim of the Lemma. We denote by $(v_i)_{i=1}^b$ the children of r, and by $(w_{i,j})_{j=1}^b$ the children of v_i. We then have

$$\begin{aligned} \mathbf{P}[\sigma_{\partial T} = h | \sigma_r = f] &= \sum_{a_1, \dots, a_b \in F} \mathbf{P}[\forall i, \sigma_{v_i}(0) = a_i | \sigma_r = f] \mathbf{P}[\sigma_{\partial T} = h | \forall i, \sigma_{v_i}(0) = a_i] \\ &= \sum_{a_1, \dots, a_b \in F} \mathbf{P}[\forall i, \sigma_{v_i}(0) = a_i | \sigma_r = f] \prod_{j=1}^b \mathbf{P}[\forall j, \sigma_{w_{i,j}} = h_{i,j} | \sigma_{v_i}(0) = a_i]. \\ &= q^{-b^2} \sum_{a_1, \dots, a_b \in F} \mathbf{P}[\forall i, \sigma_{v_i}(0) = a_i | \sigma_r = f] \prod_{j=1}^b \mathbf{P}[\forall j, \sigma_{w_{i,j}}(0) = h_{i,j}(0) | \sigma_{v_i}(0) = a_i]. \end{aligned} \tag{15}$$

Note that Lemma 3.1 implies that if $x_1, \dots, x_j$ is set of nonzero elements and $1 \leq j \leq b$, then the uniform distribution measure on $\mathcal{F}^b[x]$ satisfies for all $y_1, \dots, y_j, a \in \mathcal{F}$ that

$$\mathbf{P}[\forall\, 1 \leq i \leq j, f(x_i) = y_i | f(0) = a] = q^{-j}.$$

Thus the probability at the right hand side of (15):

$$\mathbf{P}[\forall j, \sigma_{w_{i,j}}(0) = h_{i,j}(0) | \sigma_{v_i}(0) = a_i].$$

does not depend on a_i. Now the claim follows. ■

Lemma 3.3. *Let* $\{x_1, \dots, x_{b+1}\}$ *be a fixed set of non-zero element in* $\mathcal{F}$*. For an element* $f \in \mathcal{F}^b[x]$ *satisfying* $f(x_k) = y_k$ *for* $1 \leq k \leq b+1$ *and a permutation* $\pi \in S_{b+1}$ *we define* f_π *to be the element of* $F^b[x]$ *which satisfy*

$$f_\pi(x_{\pi(k)}) = y_k. \tag{16}$$

Moreover, for $a, b \in \mathcal{F}$ *define:*

$$n_{a,b} = \left| \left\{ (f, \pi) \in F^b[x] \times S_{b+1} : f(0) = a, f_\pi(0) = b \right\} \right| . \tag{17}$$

We then have for all $a \neq b, c \neq d$ that

$$n_{a,b} = n_{c,d}, \quad n_{a,a} = n_{c,c} \ . \tag{18}$$

Moreover, for all $a \neq b$,

$$n_{a,a} > n_{a,b}. \tag{19}$$

Proof. We note that if $c \in F$, and $f \in F^b[x]$, then $(f+c)_\pi = f_\pi + c$. Similarly, if $0 \neq d \in F$, then $(df)_\pi = df_\pi$. Thus, for all a, b we have $n_{a+c,b+c} = n_{a,b}$ and $n_{da,db} = n_{a,b}$. Relations (18) now follows.

Looking at (14) one sees that the elements of $\mathcal{F}^b[x]$ which satisfy $g(0) = a$ are exactly those elements which satisfy the equation

$$\sum_{i=1}^{b+1} c_i g(x_i) = a,$$

where c_i are some none zero constants.

Thus the elements g of $\mathcal{F}^b[x]$ which satisfy $g(0) = a$ and $g_\pi(0) = b$ are exactly those elements of $\mathcal{F}^b[x]$ which satisfy the equations

$$\sum_{i=1}^{b+1} c_i g(x_i) = a, \tag{20}$$

and

$$\sum_{i=1}^{b+1} c_{\pi(i)} g(x_i) = b. \tag{21}$$

Note that by Proposition 3.1 the number of solutions of these equations, $s_\pi(a,b)$, in $\mathcal{F}^b[x]$ is the same as the number of solutions of these equation over $\mathcal{F}$ in the variables $g(x_1), \ldots, g(x_{b+1})$. It now follows that for every permutation $\pi \neq 1$ and $a \neq b$:

$$s_\pi(a,b) \leq s_\pi(a,a). \tag{22}$$

Moreover, when $\pi = 1$,

$$0 = s_\pi(a,b) < s_\pi(a,a) = q^b. \tag{23}$$

Now,

$$n_{a,b} = \sum_{\pi \in S_{b+1}} s_\pi(a,b) < \sum_{\pi \in S_{b+1}} s_\pi(a,a) = n_{a,a},$$

as needed. ■

.

Lemma 3.4. *Let $f, g \in \mathcal{F}^b[x]$. Then f and g are indistinguishable in the sense of definition 2.1 if and only if $f = g_\pi$ for some $\pi \in S_{b+1}$.*

Proof. We will write $f \sim_1 g$ to denote that f and g are indistinguishable, and $f \sim_2 g$ when there exist $\pi \in S_{b+1}$ such that $f = g_\pi$. It is clear that if $f \sim_2 g$, then $f \sim_1 g$.

On the other hand, suppose that $f \nsim_2 g$. We will show that there exist an $h \in \mathcal{F}^b[x]$ such that

$$\sum_{h' \sim_2 h} \mathbf{P}[M(f) = h'] \neq \sum_{h' \sim_2 h} \mathbf{P}[M(g) = h]. \tag{24}$$

This would imply that indeed $\sim_1 = \sim_2$ (see definition 2.1). We may write (24) equivalently,

$$\sum_{\pi \in S_{b+1}} \mathbf{P}[M(f) = h_\pi] \neq \sum_{\pi \in S_{b+1}} \mathbf{P}[M(g) = h_\pi]. \tag{25}$$

Writing $n_f(a) = |\{1 \leq j \leq b+1 : f(x_j) = a\}|$, we may write the left hand side of (25) as:

$$\frac{1}{(b+1)q^b} \sum_{\pi \in S_{b+1}} n_f(h_\pi(0)). \tag{26}$$

Summing over all $h \in \mathcal{F}^b[x]$ with $h(0) = a$ in (26) we obtain:

$$\frac{n_{a,a} n_f(a) + \sum_{b \neq a} n_{a,b} n_f(b)}{(b+1)q^b} = \frac{(n_{a,a} - n_{a,a+1}) n_f(a) + (b+1) n_{a,a+1}}{(b+1)q^b}. \tag{27}$$

(by (18) we may write $n_{a,a+1}$ for $n_{a,b}$ when $a \neq b$). Since $f \nsim_2 g$ there exists $a \in \mathcal{F}$ such that $n_f(a) > n_g(a)$. Thus:

$$\begin{aligned} \sum_{h : h(0) = a, \pi \in S_{b+1}} \mathbf{P}[M(f) = h_\pi] &= \frac{(n_{a,a} - n_{a,a+1}) n_f(a) + (b+1) n_{a,a+1}}{(b+1)q^b} \\ &> \frac{(n_{a,a} - n_{a,a+1}) n_g(a) + (b+1) n_{a,a+1}}{(b+1)q^b} = \sum_{h : h(0) = a, \pi \in S_{b+1}} \mathbf{P}[M(g) = h_\pi]. \end{aligned} \tag{28}$$

where we have used the fact that by (19) $n_{a,a} > n_{a,a+1}$. It now follows that there exists an h for which (25) holds. ■

Proof of Theorem 1.2: Fix b and let T be the 2-level b-ary tree. Let M be the channel defined at definition 3.1. Lemma 3.1 implies that for any initial distribution $I(\sigma_r, \sigma_\partial) = 0$.

On the other hand by Lemma 3.4, the channel M satisfies that f and g are indistinguishable for M if and only if $g = f_\pi$ for some π. It follows by Theorem 2.1 that when B is sufficiently large the reconstruction problem for T_B is solvable (moreover, when B is sufficiently large we may reconstruct $\{f_\pi\}_{\pi \in S_{b+1}}$ the equivalence class of f, with probability close to 1). ■

4 Improved bounds for Potts models

In this section we will focus on the following two families of channels:

- The asymmetric binary channels. These channels have the states space $\{0, 1\}$ and the matrices:
$$\mathbf{M} = \begin{pmatrix} 1-\delta_1 & \delta_1 \\ 1-\delta_2 & \delta_2 \end{pmatrix}, \tag{29}$$
with $\lambda_2(M) = \delta_2 - \delta_1$.
- The symmetric channels on q symbols. These have the state space $\{1, \dots, q\}$ and the matrices:
$$\mathbf{M} = \begin{pmatrix} 1-(q-1)\delta & \delta & \dots & \delta \\ \delta & 1-(q-1)\delta & \delta & \dots \\ \vdots & \dots & \ddots & \vdots \\ \delta & \dots & \delta & 1-(q-1)\delta \end{pmatrix}, \tag{30}$$
with $\lambda_2(M) = 1 - q\delta$.

The reconstruction problem is unsolvable for (29) when $|b\lambda_2(M)| \le 1$ and unsolvable for (30) when $0 \le b\lambda_2(M) \le 1$ (see Propositions 3 and 4 in [22]) On the other hand, Proposition 5.1 implies that when $b\lambda_2^2(M) > 1$ the reconstruction problem is solvable for these channels. Moreover, the main results of [22] states that the reconstruction problem for (29) is solvable when $b\lambda_2(M) > 1$ and δ_1 is sufficiently small; Similarly, the reconstruction problem for (30) is solvable when $b\lambda_2(M) > 1$ and q is sufficiently large. In this section we improve the existing bounds for non-reconstruction by showing that:

Proposition 4.1. *Let M be defined by the transition matrix (29). Then the reconstruction problem for T_b is unsolvable when*
$$b\frac{(\delta_2 - \delta_1)^2}{\delta_2 + \delta_1} \le 1. \tag{31}$$
Similarly, the reconstruction problem for a general tree T is unsolvable if
$$\mathrm{br}(T)\frac{(\delta_2 - \delta_1)^2}{\delta_2 + \delta_1} < 1. \tag{32}$$

Proposition 1.3. *Let M be the channel (30). Then the reconstruction problem for T_b is unsolvable when*
$$b\frac{(1-q\delta)^2}{1-(q-2)\delta} \le 1. \tag{33}$$
Similarly, the reconstruction problem for a general tree T is unsolvable if
$$\mathrm{br}(T)\frac{(1-q\delta)^2}{1-(q-2)\delta} < 1. \tag{34}$$

Propositions 1.3 and 4.1 follow from the following theorem.

Theorem 4.2. *Let* M *be a channel and let* i, j *be two states such that for all* $k \notin \{i, j\}$ *it holds that* $\mathbf{M}_{i,k} = \mathbf{M}_{j,k}$, *and there exist* $0 \le \epsilon \le 1$ *and* $\beta \ge 0, \gamma \ge 0$ *such that*

$$\begin{pmatrix} \mathbf{M}_{i,i} & \mathbf{M}_{i,j} \\ \mathbf{M}_{j,i} & \mathbf{M}_{j,j} \end{pmatrix} = \alpha \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} + \begin{pmatrix} \beta & \gamma \\ \beta & \gamma \end{pmatrix}. \tag{35}$$

If T_b *is the* b*-ary tree and* $b\alpha(1-2\epsilon)^2 \le 1$, *then* $\lim_{n\to\infty} D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) = 0$. *Similarly, if* T *is a general tree with* $\mathrm{br}(T)\alpha(1-2\epsilon)^2 < 1$, *then* $\lim_{n\to\infty} D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) = 0$.

Proof of Proposition 4.1: If $\delta_1 + \delta_2 \le 1$ write:

$$\begin{pmatrix} 1-\delta_1 & \delta_1 \\ 1-\delta_2 & \delta_2 \end{pmatrix} = (\delta_1+\delta_2) \begin{pmatrix} \frac{\delta_2}{\delta_1+\delta_2} & \frac{\delta_1}{\delta_1+\delta_2} \\ \frac{\delta_1}{\delta_1+\delta_2} & \frac{\delta_2}{\delta_1+\delta_2} \end{pmatrix} + \begin{pmatrix} 1-\delta_1-\delta_2 & 0 \\ 1-\delta_1-\delta_2 & 0 \end{pmatrix}$$

and the proposition follows from Theorem 4.2. Otherwise, write

$$\begin{pmatrix} 1-\delta_1 & \delta_1 \\ 1-\delta_2 & \delta_2 \end{pmatrix} = \begin{pmatrix} \delta_1' & 1-\delta_1' \\ \delta_2' & 1-\delta_2' \end{pmatrix}$$

where $\delta_1' + \delta_2' < 1$, and use a similar decomposition. ■

Proof of Proposition 1.3: For all $i \ne j$ we have

$$\begin{pmatrix} \mathbf{M}_{i,i} & \mathbf{M}_{i,j} \\ \mathbf{M}_{j,i} & \mathbf{M}_{j,j} \end{pmatrix} = (1-(q-2)\delta) \begin{pmatrix} \frac{1-(q-1)\delta}{1-(q-2)\delta} & \frac{\delta}{1-(q-2)\delta} \\ \frac{\delta}{1-(q-2)\delta} & \frac{1-(q-1)\delta}{1-(q-2)\delta} \end{pmatrix}.$$

and the proposition follows from Theorem 4.2. ■

Proof of Theorem 4.2: Assume that the label of the root is chosen to be i or j with probability 1/2 each and let X denote this random label. For a set W, we denote by $\widetilde{Y}_W$ the random labeling of the vertices of W, and $\widetilde{Y}_n = \widetilde{Y}_{L_n}$. We will show that

$$\lim_{n\to\infty} I(X, \widetilde{Y}_n) = 0, \tag{36}$$

where I is the mutual information operator, i.e., $I(X,Y) = H(X)+H(Y)-H(X,Y)$ where H is the entropy operator (see [7] for more background). Equation (36) is equivalent to $\lim_{n\to\infty} D_V(\mathbf{P}_n^i, \mathbf{P}_n^j) = 0$, see e.g. [22].

We will split the proof into three main steps.

Step 1 [10]: If M is the binary symmetric channel,

$$\mathbf{M} = \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix}, \tag{37}$$

then for a general tree T and a set W,

$$I(X, \widetilde{Y}_W) \le \sum_{w\in W} (1-2\epsilon)^{2|w|}. \tag{38}$$

This is Theorem 1.3 in [10].

Step 2: We will show that it suffices to prove the theorem assuming that

$$\begin{pmatrix} \mathbf{M}_{i,i} & \mathbf{M}_{i,j} \\ \mathbf{M}_{j,i} & \mathbf{M}_{j,j} \end{pmatrix} = \alpha \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} \tag{39}$$

instead of (35). Indeed, assume that the theorem is true under the condition (39). Now let M be a channel which satisfies (35). Consider the following auxiliary channel N. The channel has the state space $\mathcal{A}' = \mathcal{A} \cup \{i^*, j^*\}$ and the following transition matrix $\mathbf{N}$:

1. For all $\ell \notin \{i,j\}$ and all $\ell' \in \mathcal{A}$, set $\mathbf{N}_{\ell,\ell'} = \mathbf{M}_{\ell,\ell'}$.
2. $$\begin{pmatrix} \mathbf{N}_{i,i} & \mathbf{N}_{i,j} \\ \mathbf{N}_{j,i} & \mathbf{N}_{j,j} \end{pmatrix} = \alpha \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix}.$$
3. $$\begin{pmatrix} \mathbf{N}_{i,i^*} & \mathbf{N}_{i,j^*} \\ \mathbf{N}_{j,i^*} & \mathbf{N}_{j,j^*} \end{pmatrix} = \begin{pmatrix} \beta & \gamma \\ \beta & \gamma \end{pmatrix}.$$
4. For all ℓ, set $\mathbf{N}_{i^*,\ell} = \mathbf{N}_{i,\ell}$ and $\mathbf{N}_{j^*,\ell} = \mathbf{N}_{j,\ell}$.

It is clear that if the original channel M satisfied the conditions of Theorem 4.2, then the channel N satisfies these conditions with (39) replacing (35). By our assumption this implies that $\lim_{n\to\infty} I(X, \widehat{Y}_n) = 0$, where $\widehat{Y}_n$ is the labeling of level n for N.

Writing $\widetilde{Y}_n$ for the labeling of level n for M, we note that $\widetilde{Y}_n$ may be obtained from $\widehat{Y}_n$ by replacing each occurrence of i^* by i and each occurrence of j^* by j. By the Data Processing Lemma (see e.g. [7]) it now follows that

$$I(X, \widetilde{Y}_n) \le I(X, \widehat{Y}_n) \to 0,$$

as needed.

Step 3: We the theorem assuming that (39) holds. We begin by introducing two random variables Z and Y. Recall that $\texttt{path}(v)$ is the path from ρ to v. We let:

$$Z = \{(v, \sigma_v) : \exists w \in \texttt{path}(v), \sigma_w \notin \{i,j\}\}, \tag{40}$$

and given a set W, let

$$Y_W = \{(v, \sigma_v) : v \in W, \forall w \in \texttt{path}(v), \sigma_w \in \{i,j\}\}. \tag{41}$$

We denote Y_{L_n} by Y_n. Roughly speaking, Y_n contains all information on the $i \to j$ and $j \to i$ process from the root to level n; Z contains all information which is independent of this process.

By the Data Processing Lemma

$$I(X, \widetilde{Y}_n) \le I(X, (Y_n, Z)). \tag{42}$$

Since for all $\ell \notin \{i, j\}$ it holds that $\mathbf{M}_{i,\ell} = \mathbf{M}_{j,\ell}$, it follows that Z is independent of X. We therefore obtain

$$\begin{aligned} I(X, (Y_n, Z)) &= H(X) + H(Y_n, Z) - H(X, Y_n, Z) = H(X|Z) + H(Y_n|Z) - H(X, Y_n|Z) \\ &= \mathbf{E}_z I(X, Y_n | Z = z). \end{aligned} \tag{43}$$

We will show that for almost all z,

$$\lim_{n \to \infty} I(X, Y_n | Z = z) = 0. \tag{44}$$

Since $0 \le I(X, Y_n | Z = z) \le 1$ for all z, this would imply the theorem by (43) and (42). By the Data Processing Lemma, in order to prove (44), it suffices to prove that for a.e. z there exist cutsets W_n for which

$$\lim_{n \to \infty} I(X, Y_{W_n} | Z = z) = 0. \tag{45}$$

(Recall that W is a cutset if W intersects every infinite path emanating from ρ; If W_n is a cutest, then and if for all $v \in W_n$ we have $\ell \le |v| \le \ell'$, then by the Data Processing Lemma it follows

$$I(X, Y_{\ell'}) \le I(X, Y_{W_n}) \le I(X, Y_\ell),$$

so (44) is equivalent to (45)).

The key observation is noting that given $Z = z$ we have a broadcast process with the binary symmetric channel (37) on the tree

$$T_z = \{v : \forall w \in \mathtt{path}(v), \sigma_w \in \{i, j\}\}.$$

Therefore by Step 1, in order to prove (45) it suffices to show that for a.e. z, there exist cutsets W_n such that

$$\lim_{n \to \infty} \sum_{w \in W_n} (1 - 2\epsilon)^{2|w|} = 0. \tag{46}$$

However, for a general tree T, we have $\mathrm{br}(T_z) = \alpha \mathrm{br}(T)$ for a.e. z. So if $\alpha \mathrm{br}(T)(1 - 2\epsilon)^2 < 1$, then (46) hold for a.e. z for appropriate W_n's. Equation (46) holds for T_d when $b\alpha(1 - 2\epsilon)^2 = 1$ by [19, Theorem 3]. ■

5 $\mathrm{br}(T)|\lambda_2|^2 > 1$ implies reconstruction

The following proposition is a consequence of the results of [17].

Proposition 5.1. *Suppose that $b|\lambda_2(M)|^2 > 1$, then the reconstruction problem is census solvable.*

In this section we prove an extension of the proposition to general trees.

Theorem 1.4. *Let T be an infinite tree and M a channel such that* $\mathrm{br}(T)|\lambda_2(M)|^2 > 1$, *then the reconstruction problem for M on T is solvable.*

In order to prove this theorem we are going to bound from below $D_{\chi^2}(\mathbf{P}_n^{(c),i}, \mathbf{P}_n^{(c),j})$ where

$$D_{\chi^2}(P, Q) = \frac{1}{2}\sqrt{\sum_{\sigma} \frac{2(P(\sigma) - Q(\sigma))^2}{P(\sigma) + Q(\sigma)}}, \tag{47}$$

is the χ^2 distance.

Lemma 5.2.

$$D_{\chi^2}^2 \leq D_V \leq D_{\chi^2}. \tag{48}$$

Proof. The first inequality in (48) follows when we use the estimate:

$$\frac{2|P(\sigma) - Q(\sigma)|}{P(\sigma) + Q(\sigma)} \leq 2,$$

while the second follows from Cauchy-Schwartz when we write:

$$\frac{1}{2}|P(\sigma) - Q(\sigma)| = \frac{|P(\sigma) - Q(\sigma)|}{\sqrt{2(P(\sigma) + Q(\sigma))}}\sqrt{\frac{P(\sigma) + Q(\sigma)}{2}}.$$

■

We prove Theorem 1.4 by constructing linear estimators of the root variable for finite trees and then evaluating the first and second moments of these estimators.

Abbreviate $\lambda = \lambda_2(M)$, and take v to be a right eigenvector (that is, $\mathbf{M}v = \lambda v$) with $|v|_2 = 1$. We use the notion of *flows* and view the tree as an *electrical network*. We refer the reader to [8] and [10] for definitions and more background. For a finite tree T, we consider T as an electrical network, where we assign the edge e the resistance

$$R(e) = (1 - |\lambda|^2)|\lambda|^{-2|e|}.$$

For a vertex x, we let $c_x = e_i$ if the label of x is i, i.e. if $\sigma_x = i$.

We say that a set of vertices W is an *anti-chain* if no vertex in W is a descendant of another.

Lemma 5.3. *Let T be a finite tree, W an anti-chain in T, and μ a unit flow from ρ to W. Consider the estimator*

$$S_\mu = \sum_{x \in W} \frac{\mu(x) c_x \cdot v}{\lambda^{|x|}}. \tag{49}$$

Then for all $\ell \in \mathcal{A}$,

$$\mathbf{E}^\ell[S_\mu] = e_\ell \cdot v, \tag{50}$$

and there exist a constant $0 < c(M) < 1$ which depends on M only such that

$$c(M)(1 + \mathcal{R}_{\text{eff}}(\rho \leftrightarrow W)) \leq \min_\mu \min_l \mathbf{E}^\ell[|S_\mu|^2] \leq \min_\mu \max_l \mathbf{E}^\ell[|S_\mu|^2] \leq (1 + \mathcal{R}_{\text{eff}}(\rho \leftrightarrow W)). \tag{51}$$

Proof. If x is at level n, then

$$\mathbf{E}^i_n[c_x] = e_i\mathbf{M}^n. \tag{52}$$

It follows from (52) that for every column vector v we have

$$\mathbf{E}^i_n[c_x \cdot v] = e_i\mathbf{M}^n v. \tag{53}$$

If we take v to be a vector such that v is an eigenvalue of $\mathbf{M}$ which corresponds to $\lambda = \lambda_2(M)$, we obtain,

$$\mathbf{E}^\ell_n[c_x \cdot v] = \lambda^n e_\ell \cdot v. \tag{54}$$

Now (50) follows by linearity.

We let $\mathtt{path}(x)$ be the path from ρ to x (more generally let $\mathtt{path}(x, y)$ be the path from x to y). We let $x \wedge y$, be the *meeting point* of x and y, that is, the vertex farthest from the root ρ on $\mathtt{path}(x) \cap \mathtt{path}(y)$. We have:

$$\begin{aligned} \mathbf{E}^\ell[|(c_x \cdot v)(c_y \cdot v)|] &= \sum_i \mathbf{P}^\ell[c_{x\wedge y} = e_i]\mathbf{E}[|(c_x \cdot v)(c_y \cdot v)|\,|c_{x\wedge y} = e_i] \\ &= \sum_i \mathbf{P}^\ell[c_{x\wedge y} = e_i]|e_i \cdot v|^2|\lambda|^{d(x,y)} \le |\lambda|^{d(x,y)}. \end{aligned} \tag{55}$$

Therefore,

$$\mathbf{E}^\ell[|S_\mu|^2] \le \sum_{x,y} \frac{\mu(x)\mu(y)}{|\lambda|^{|x|}|\lambda|^{|y|}}\mathbf{E}^\ell[|(c_x \cdot v)(c_y \cdot v)|] \le \sum_{x,y} \frac{\mu(x)\mu(y)}{|\lambda|^{2|x\wedge y|}}. \tag{56}$$

Since

$$\frac{1}{|\lambda|^{2|u|}} = 1 + \sum_{e\in\mathtt{path}(u)} R(e),$$

if follows that for all ℓ,

$$\mathbf{E}^\ell[|S_\mu|^2] \le (1 + \sum_e R(e) \sum_{x,y\in W} \mathbf{1}_{e\in\mathtt{path}(x\wedge y)}\mu(x)\mu(y)) = (1 + \sum_e R(e)\mu^2(e)). \tag{57}$$

When we take the minimum in (57) we obtain,

$$\min_\mu \max_\ell \mathbf{E}^\ell[|S_\mu|^2] \le \min_\mu(1 + \sum_e R(e)\mu^2(e)) = (1 + \mathcal{R}_{\text{eff}}(\rho \leftrightarrow W))\,, \tag{58}$$

where the equality in (58) is Thompson's principle (see Doyle and Snell [8]). Equation (58) is the left hand side inequality in (51). The proof of the other inequality in (51) follows in a similar way. ■

Proof of Theorem 1.4: Since $|\lambda|^2 \mathrm{br}(T) > 1$ it follows that

$$\mathcal{R}_{\mathrm{eff}}(\rho \leftrightarrow \infty) = \sup_n \mathcal{R}_{\mathrm{eff}}(\rho \leftrightarrow L_n) < \infty. \tag{59}$$

We consider linear estimators as in Lemma 5.3 and note that if v is an eigenvector as in Lemma 5.3, then

$$\widetilde{C}(M) = \max_{i,j} |e_i \cdot v - e_j \cdot v| > 0. \tag{60}$$

We fix a level n, take i and j such that we obtain the maximum in (60), and consider a linear estimator as in Lemma 5.3 for $W = L_n$, such that the upper bound in (51) holds. We obtain that

$$\left| \sum_\sigma S_\mu(\sigma)(\mathbf{P}_n^i[\sigma] - \mathbf{P}_n^j[\sigma]) \right| \geq \widetilde{C}(M). \tag{61}$$

On the other hand, by Cauchy-Schwartz we obtain that

$$\begin{aligned} \left(\sum_\sigma S_\mu(\sigma)(\mathbf{P}_n^i[\sigma] - \mathbf{P}_n^j[\sigma]) \right)^2 &\leq \sum_\sigma |S_\mu(\sigma)|^2 (\mathbf{P}_n^i[\sigma] + \mathbf{P}_n^j[\sigma]) \sum_\sigma \frac{(\mathbf{P}_n^i[\sigma] - \mathbf{P}_n^j[\sigma])^2}{\mathbf{P}_n^i[\sigma] + \mathbf{P}_n^j[\sigma]} \\ &\leq 4(1 + \mathcal{R}_{\mathrm{eff}}(\rho \leftrightarrow L_n)) D^2_{\chi^2}(\mathbf{P}_n^i, \mathbf{P}_n^j). \end{aligned} \tag{62}$$

Combining (61) and (62) we obtain that for all n

$$D^2_{\chi^2}(\mathbf{P}_n^i, \mathbf{P}_n^j) \geq \frac{\widetilde{C}(M)^2}{4(1 + \mathcal{R}_{\mathrm{eff}}(\rho \leftrightarrow L_n))}. \tag{63}$$

So the theorem follows by Lemma 5.2. ■

Proof of Proposition 5.1: For the b-ary tree T_b we have $\mathrm{br}(T_b) = b$. Moreover, the flow μ which minimizes $\sum_e R(e)\mu^2(e)$ in (58) is the flow which satisfies $\mu(e) = b^{-|e|}$. This implies that we may take S_n to be

$$S_n = \frac{c_n \cdot v}{d^n \lambda^n}.$$

In particular S_n is a function of c_n. Now, arguing as in the proof of Theorem 1.4 we see that there exist i, j for which $D_V(\mathbf{P}_n^{(c),i}, \mathbf{P}_n^{(c),j})$ is bounded away from 0 for all n. ■

6 Census reconstruction fails when $b|\lambda_2|^2 < 1$

In this section we prove non-reconstruction when $b|\lambda_2|^2 < 1$. The proof relies on the Kesten-Stigum theorem [17]. The following is an immediate consequence of Theorem 2.3 of [17].

Kesten-Stigum CLT [17]: *Let M be a transition matrix such that $b|\lambda_2(M)|^2 < 1$ and let π be the stationary distribution for M, i.e., the normalized left eigenvector for the eigenvalue 1. Then for any vector v of M which is orthogonal to π, and for all $j \in \mathcal{A}$,*

$$\frac{(c_n^j, v)}{b^{n/2}} \to \mathcal{N}(0, \sigma), \tag{64}$$

where $\mathcal{N}$ denotes a normal random variable and σ does not depend on j. Moreover, the convergence rate may be bounded in term of $|v|_2$ only.

An immediate consequence is:

Proposition 6.1. *For all i, j and $\epsilon > 0$, one may couple c_n^j and c_n^i in such a way that,*

$$\lim_{n\to\infty} \mathbf{P}[\,|\,c_n^j - c_n^i\,| > \epsilon b^{n/2}] = 0. \tag{65}$$

Proof. If follows from the Kesten Stigum theorem that

$$\frac{c_n^\ell - \pi b^n}{b^{n/2}} \to \mathcal{N}$$

where $\mathcal{N}$ is a normal variable which does not depend on ℓ. The claim follows, as for all $\delta > 0$ we can couple both c_n^i and c_n^j with the same normal variable $b^{n/2}\mathcal{N}$ in such a way that

$$\mathbf{P}[\,|\,c_n^\ell - b^{n/2}\mathcal{N}\,| > \epsilon b^{n/2}] < \delta,$$

for both $\ell = i$ and $\ell = j$. ■

Theorem 1.5. *If $b|\lambda_2(M)|^2 < 1$, then the reconstruction problem is not census solvable for the b-ary tree and the channel M.*

The proof is easier when all the entries of $\mathbf{M}$ are strictly positive, in which case the following lemma is trivial.

Lemma 6.2. *Let M be irreducible and aperiodic, then there exists an h such that for all i and j all $n \geq h$ and all v,*

$$\mathbf{P}[c_n^i = v] > 0 \textit{ iff } \mathbf{P}[c_n^j = v] > 0. \tag{66}$$

Proof. Since $(\mathbf{M}_{i,j})_{i,j=1}^k$ is irreducible and aperiodic it follows that there exists an ℓ such that the matrix $\mathbf{M}^\ell$ is strictly positive. Thus, for every two states i and j we may construct a ℓ level tree such that the root of the tree is labeled by i and all of the leaves are labeled by j. Let $h = \ell(k+4)$ and $n \geq h$. Suppose that $\mathbf{P}[c_n^i = v] > 0$. Then there exists a labeling of the tree of n levels which has i at the root and σ at level n, and such the census of σ is v. We will prove the lemma by constructing a labeling where the root is labeled by j, level n is labeled by τ, and the census of τ is v.

We denote by $x_1, \ldots, x_{b^{n-\ell}}$ the vertices at level $n - \ell$. For x_t we define $c(x_t)$ to be the census of the subtree rooted at x_t for the labeling σ. Let $u = (u_1, \ldots, u_k)$ be a non-negative vector which satisfies $\sum u_i = b^\ell$. Define $c(\sigma, u) = |\{t : c(x_t) = u\}|$. Note that if $c(\sigma, u) > 0$ then there exists a labeling of the ℓ levels tree, denoted $\tau(u)$, which has u as its census and such that the label of the root for this labeling is $\sigma_\rho(u)$.

Since

$$v - \sum_u \left(c(\sigma, u) \mod b^\ell\right) u = 0 \mod b^\ell,$$

it follows that

$$v = b^\ell w + \sum_u \left(c(\sigma, u) \mod b^\ell\right) u,$$

for some integer valued vector w.

Note that if for root value j, we could label $\sum_u \left(c(\sigma, u) \mod b^\ell\right)$ of the vertices of level $n-\ell$ by labels $\sigma_\rho(u)$ with multiplicity $c(\sigma, u) \mod b^\ell$, then we are done. Since using these vertices, it is possible to build the $\sum_u \left(c(\sigma, u) \mod b^\ell\right) u$ part of the census. Now using the other vertices at level $n-\ell$ it is possible to build the $b^\ell w$ part of the census (whatever labels these vertices have). Note that

$$\sum_u \left(c(\sigma, u) \mod b^\ell\right) \leq b^\ell |\{u : \forall i, u_i \geq 0, \sum_{i=1}^k u_i = b^\ell\}| \leq b^\ell b^{k\ell} = b^{(k+1)\ell}.$$

Therefore by assigning $b^{(k+1)\ell}$ of the vertices at level $n-2\ell$ the task of producing the prescribed labels at level $n-\ell$ we obtain the required result. ■

Proof of Theorem 1.5: We take h such that (66) holds for $n \geq h$. We may write:

$$c^j_{n+h} = \sum_{i=1}^k S^i(c^j_n(i)), \tag{67}$$

where S^i is a random walk on $\mathbf{Z}^k$ which satisfy:

$$\mathbf{P}[S^i(t+1) = S^i(t) + v] = \mathbf{P}[c^i_h = v].$$

By (66), all the S^i are random walks with the same support.

By the local CLT if follows that if $|c^j_n - c^i_n| < \epsilon\sqrt{b^n}$, it is possible to couple c^j_{n+h} and c^i_{n+h} with probability $f(\epsilon)$ where $f(\epsilon) \to 0$ as $\epsilon \to 0$. Now the result follows from Proposition 6.1.

■

7 Census reconstruction for Potts models fails at criticality

In section 6 we presented a proof that $b|\lambda_2|^2 < 1$ implies that the reconstruction problem is not census-solvable. We believe that this result is also valid when $b|\lambda_2|^2 = 1$. In this section we prove that this is the case when M is the q state Potts model or the Ising model with external field.

Theorem 1.6. *Let M be the q state Potts model, or the asymmetric Ising model, and suppose that $b\lambda_2^2(M) \leq 1$, then the reconstruction problem is unsolvable.*

Proof. We prove the theorem for the 3-state Potts model, the general proof being similar. It is convenient to denote the states of the channel by $-1, 0$ and 1. We show that for all $i, j \in \{-1, 0, 1\}$ there exist a coupling such that $\mathbf{P}[c^i_n = c^j_n] \to 1$ as $n \to \infty$. By symmetry, it suffices to find such coupling for c_n^1 and c_n^{-1}. We denote by σ^+ the coloring with 1 at the root, and σ^- for the coloring with -1 at the root.

During all the steps of the coupling we require that for all v we have $\sigma_v^- = 0$ iff $\sigma_v^+ = 0$. The main step in proving the existence of the required coupling is given in the following lemma.

Lemma 7.1. *There exits $p^* > 0$ such that: Given $(\sigma_v)_{|v| \le n}$ and $(\tau_v)_{|v| \le n}$ which are coupled in such a way that $\sigma_v = 0$ iff $\tau_v = 0$, there exits $N \ge n$ and a coupling procedure for $(\sigma_v)_{|v| \le N}$ and $(\tau_v)_{|v| \le N}$ such that $\sigma_v = 0$ iff $\tau_v = 0$ and such that $\mathbf{P}[c_N^\sigma = c_N^\tau] \ge p^*$ (where c_N^σ and c_N^τ are the σ and τ level N census vectors respectively).*

By iterating Lemma 7.1, the coupling probability $\mathbf{P}[c_n^1 = c_n^{-1}]$ may be made arbitrarily close to 1 for large n. ∎

Proof of Lemma 7.1: At level n there are two types of vertices: those for which $\sigma_v = \tau_v$, and those for which $\sigma_v = -\tau_v \in \{-1, 1\}$. For all the vertices w whose ancestors v at level n satisfy $\sigma_v = \tau_v$, we let $\sigma_w = \tau_w$.

Denote the vertices v at level n which satisfy $\sigma_v = -\tau_v \in \{-1, 1\}$ by $v_1, \ldots, v_\ell$. Let T be the graph which consists of the subtrees rooted at $v_1, \ldots, v_\ell$.

We think of T as drawn in the plan, and denote the vertices of level m by $v_1^m, \ldots, v_{\ell b^{n-m}}^m$. We will slightly abuse the notation by redefining

$$c_m^\sigma(i) = |\{j : 1 \le j \le \ell b^{n-m}, \sigma(v_j^m) = i\}|, \quad c_m^\tau(i) = |\{j : 1 \le j \le \ell b^{n-m}, \tau(v_j^m) = i\}|. \tag{68}$$

Note that $c_m^\sigma = c_m^\tau$ in the old definition if and only if it does in the new one.

We define $X_m^\sigma = c_m^\sigma(1) - c_m^\sigma(-1)$ and $X_m^\tau = c_m^\tau(1) - c_m^\tau(-1)$. Note that if a coupling of σ and τ satisfies $\sigma_v = 0$ iff $\tau_v = 0$, then we have $c_m^\sigma = c_m^\tau$ if and only if $X_m^\sigma = X_m^\tau$.

We define the coloring of T as a dynamic process starting at the root, running level by level from left (v_1^m) to right $(v_{\ell b^{n-m}}^m)$. Writing v' for the parent of v, we let

$$S^\sigma(v_i^m) = (b\theta)^{-n} X_n^\sigma + \sum_{n \le k < m} (b\theta)^{-k} \sum_{j \le \ell b^{k-n}} (\sigma(v_j^k) - \theta\sigma(v_j^{k'})) + (b\theta)^{-m} \sum_{j \le i} (\sigma(v_j^m) - \theta\sigma(v_j^{m'})), \tag{69}$$

where we formally define $\sigma_{\rho'} = 0$. We define S^τ similarly. Note that all the terms in (69) except the first one, have mean zero, and therefore both S^σ and S^τ are martingales. Also,

$$S^\sigma(v_{\ell b^{m-n}}^m) = (b\theta)^{-m} X_m^\sigma, \quad S^\tau(v_{\ell b^{m-n}}^m) = (b\theta)^{-m} X_m^\tau. \tag{70}$$

The coupling consists of a few steps, during all of which we require that for all v we have $\sigma_v = 0$ iff $\tau_v = 0$.

1. **Reflection until we reach N such that:**

$$-\sqrt{b^N} C_1 \le X_N^\sigma = -X_N^\tau \le C_1 \sqrt{b^N}. \tag{71}$$

We label the vertices level by level, left to right using the rule $\sigma_v = -\tau_v$, until we reach v_i^N with $|S^\sigma(v_i^N)| \le (1+\theta)(b\theta)^{-N}$ and N large. From lemma 7.2 below, it follows that such a vertex exits a.s. Now $X_N^\sigma - (b\theta)^N S^\sigma(v_i^N)$ is a sum of $\ell b^{N-n} - i$ independent variables with values $-1, 0, 1$ and expected value 0. It therefore follows by the CLT that with probability going to 1 as $C_1 \to \infty$ we have $|X_N^\sigma| \le C_1\sqrt{b^N}$. Therefore, by continuing the reflection $\sigma(v_j^N) = -\tau(v_j^N)$ for $i < j \le \ell b^{N-n}$, we obtain (71). We let $Y_{N+1}^\sigma(i,j)$ be the number of vertices in the σ labeling which are at level $N+1$, of type j and have a type i parent. All these variables are binomial and

$$\mathbf{E}[Y_{N+1}^{\sigma,\tau}(i,j)|c_N^{\sigma,\tau}] = (b(1-\theta)/3 + \delta_{i,j}b\theta)\, c_N^{\sigma,\tau}(i). \tag{72}$$

2. **Labeling the 0's of level $N+1$:** We label all vertices of level $N+1$ with $\sigma(v_j^{N+1}) = \sigma(v_j^{N+1}) = 0$ independently with probability

   $$\mathbf{P}[\sigma(v_j^{N+1}) = 0|\sigma((v_j^{N+1})')] = \mathbf{P}[\tau(v_j^{N+1}) = 0|\tau((v_j^{N+1})')].$$

   We thus achieve that $c_{N+1}^\sigma(0) = c_{N+1}^\tau(0)$.

   By the CLT it follows that with probability going to 1 as $C_3 \to \infty$,

   $$|Y_{N+1}^{\sigma,\tau}(i,0) - \mathbf{E}[Y_{N+1}^{\sigma,\tau}(i,0)|c_N^{\sigma,\tau}]| \le C_3\sqrt{b^{N+1}},$$

   and therefore for $j \ne 0$,

   $$|\mathbf{E}[Y_{N+1}^{\sigma,\tau}(i,j)|c_N^{\sigma,\tau}, Y_{N+1}^{\sigma,\tau}(i,0)] - \mathbf{E}[Y_{N+1}^{\sigma,\tau}(i,j)|c_N^{\sigma,\tau}]| \le C_3\sqrt{b^{N+1}}. \tag{73}$$

3. **Labeling some of the $\pm$ of level $N+1$.** Let i_0 be the maximizer of $c_N^\sigma(i)$. For $i \ne i_0$ we label all vertices v at level $N+1$ which have a parent of type i and satisfy $\sigma(v) \ne 0$ by ± 1 using the reflecting coupling $\sigma(v) = -\tau(v)$. By the CLT it follows that with probability going to 1 as $C_4 \to \infty$ we have for $i \ne i_0$ and all j

   $$|Y_{N+1}^{\sigma,\tau}(i,j) - \mathbf{E}[Y_{N+1}^{\sigma,\tau}(i,j)|c_N^{\sigma,\tau}, Y_{N+1}^{\sigma,\tau}(i,0)| \le C_4\sqrt{b^{N+1}}. \tag{74}$$

4. **Local CLT for the final coupling.** We now want to label the remaining vertices with ± 1 to achieve $X_{N+1}^\sigma = X_{N+1}^\tau$. We note that $Z^{\sigma,\tau} = Y^{\sigma,\tau}(i_0,1) - Y^{\sigma,\tau}(i_0,-1)$ are sum of i.i.d. ± 1 random variables and therefore satisfy the local CLT. Moreover, by (73), it follows that with high probability each of the Z's is a sum of at least $\ell b^{N+1-n}(1-\theta)/6$ such variables. In order that $X_{N+1}^\sigma = X_{N+1}^\tau$ we need that $\sum_i Y^\sigma(i,1) - Y^\sigma(i,-1) = \sum_i Y^\tau(i,1) - Y^\tau(i,-1)$. By (72),(73) and (74) it follows that for $i \ne i_0$ and with probability going to 1 as $C_5 \to \infty$,

   $$|(Y^{\sigma,\tau}(i,1) - Y^{\sigma,\tau}(i,-1)) - b\theta c_N^{\sigma,\tau}(i)| \le C_5\sqrt{b^{N+1}}. \tag{75}$$

   Similarly, by (72) and (73)

   $$|\mathbf{E}[Y^{\sigma,\tau}(i_0,1) - Y^{\sigma,\tau}(i_0,-1)|c_N^{\sigma,\tau}, Y_{N+1}^{\sigma,\tau}(i,0)] - b\theta c_N^{\sigma,\tau}(i_0)| \le C_5\sqrt{b^{N+1}} \tag{76}$$

   Now by (75),(76),(71) and the fact that until now we used reflection, it follows that $Z^\sigma + W^\sigma = Z^\tau + W^\tau$ where $W^{\sigma,\tau}$ are random variables which satisfies $W^\sigma = -W^\tau$ and with probability going to 1 as $C_6 \to \infty$, $|W^\sigma - W^\tau - (\mathbf{E}[Z^\sigma] - \mathbf{E}[Z^\tau])| \le C_6\sqrt{b^{N+1}}$. By the local CLT, it follows the with probability $p^* > 0$ we may couple Z^σ and Z^τ in order to achieve $X_N^\sigma = X_N^\tau$.

Lemma 7.2. *A.s. S^σ changes signs infinitely often.*

Proof. We prove that X_n^σ changes sign infinitely often a.s. By the Borel-Cantelli lemma it suffices to prove that for all n and all $\sigma(v)_{|v|\le n}$ there exists $m \ge n$ such that $\mathbf{P}[X_m^\sigma > 0|\sigma(v)_{|v|\le n}] \ge 1/4$. However, by the Kesten-Stigum theorem [17], there exits some $a_m \to \infty$ such that given $\sigma(v)_{|v|\le n}$, X_m/a_m converges to a non-degenerate normal random variable with expected value 0. So when m is large, $\mathbf{P}[X_m^\sigma > 0|\sigma(v)_{|v|\le n}] \ge 1/4$ as needed. ■

8 Unsolved problems

8.1 Critical values

Except for symmetric binary channels there are no interesting families of channels for which the critical value for reconstruction is known.

Problem 1. *For the* 3 *symbols Potts model (30) find the critical value for reconstruction.*

The same question applies to asymmetric binary channels, and to proper colorings (which correspond to the zero temperature anti-ferromagnetic Potts model). Proper colorings of trees were studied in [2]; The corresponding transition matrix is

$$\mathbf{M} = \begin{pmatrix} 0 & (q-1)^{-1} & (q-1)^{-1} & \dots & (q-1)^{-1} \\ (q-1)^{-1} & 0 & (q-1)^{-1} & \dots & \\ \vdots & \dots & \ddots & \vdots & \\ (q-1)^{-1} & \dots & (q-1)^{-1} & 0 & \end{pmatrix}. \tag{77}$$

A simple coupling argument (see [2]) shows that if $b \le q-1$, then the reconstruction problem is unsolvable for colorings of T_b. On the other hand, applying standard coupon-collector estimates recursively (similarly to Theorem 2.1), it is easy to see that if $b \ge (1+\epsilon)q\log q$ and q is large, then the reconstruction problem is solvable. By Proposition 5.1 and Theorem 1.5, the census reconstruction problem is solvable if $b > (q-1)^2$, and unsolvable if $b < (q-1)^2$.

Problem 2. *For fixed q, for which b is the reconstruction problem solvable for the channel (77)?*

8.2 Soft inputs - Robust phase transitions

We like to mention briefly the notion of "robust" phase transition which first appeared in [23]. Consider the usual reconstruction problem, but suppose that the data at the boundary is given with some additional noise. The proofs that if $b\lambda_2^2(M) > 1$ the reconstruction problem is (census) solvable are immune to this noise. However, this may not be the case for the reconstruction problem. Indeed, we suspect that adding this additional noise (assuming it is fixed but sufficiently strong) will shift the phase transition to the point $b\lambda_2^2(M) = 1$. A similar phenomena was proven in [23] for the phase transition of uniqueness. For n and m, we denote by $\sigma_{n,m}$ the configuration which is obtained from σ_n by applying the random function M^m independently on each of the symbols in σ_n. We denote by $\mathbf{P}_{n,m}^\ell$ the conditional distribution on $\sigma_{n,m}$ given that $\sigma_\rho = \ell$. We then

Conjecture 1. *For all* M *and* b*, such that* $b\lambda_2^2(M) < 1$*, there exists* m *such that*

$$\sup_{i,j} \lim_{n\to\infty} D_V(\mathbf{P}^i_{n,m}, \mathbf{P}^j_{n,m}) = 0.$$

8.3 Monotonicity

For Potts models (4), it is easy to see that if the reconstruction problem is solvable for q and ϵ and $q' < q$, then the reconstruction problem is also solvable for q' and ϵ. We expect that for fixed $\lambda = \lambda_2(M)$, reconstruction is easier when q is larger.

Conjecture 2. *Consider two symmetric channels* M_1 *and* M_2 *(as in (4)) on* q_1 *and* q_2 *symbols respectively, where* $q_1 < q_2$*. If* $\lambda_2(M_1) = \lambda_2(M_2)$ *and the reconstruction problem is solvable for* M_1*, then it is also solvable for* M_2*.*

This is obvious when q_2 is a multiple of q_1. Using the reconstruction criteria for the binary symmetric channel on 2 symbols, it is easy to prove the conjecture when $q_1 = 2$.

Acknowledgments: We thank Olle Häggström, Claire Kenyon, Làszlò Lovàsz, Jeff Steif and Peter Winkler for helpful discussions.

References

[1] Athreya, K. B. and Ney, P. E. (1972) *Branching Processes*, Springer-Verlag.

[2] Brightwell, G. and Winkler, P. (2001). Random colorings of a Cayley tree, *submitted.*

[3] Brightwell, G. and Winkler, P. (2000). Gibbs measures and dismantlable graphs, *J. Comb. Theory (Series B)* **78**, 141–169.

[4] Brightwell, G. and Winkler, P. (1999). Graph homomorphisms and phase transitions, *J. Comb. Theory (Series B)*, **77**, 415–435.

[5] Bleher, P. M., Ruiz, J. and Zagrebnov V. A. (1995) On the purity of limiting Gibbs state for the Ising model on the Bethe lattice, *J. Stat. Phys* **79**, 473–482.

[6] Cavender, J. (1978). Taxonomy with confidence. *Math. BioSci.* **40**, 271–280.

[7] Cover, T. M. and Thomas, J. A. (1991) *Elements of Information Theory*, John Wiley and Sons.

[8] Doyle, P. G. and Snell, E. J.(1984). *Random walks and Electrical Networks.* Carus Math. Monographs **22**, Math. Assoc. Amer., Washington, D. C.

[9] Dembo, A. and Zeitouni O. (1997) *Large Deviations, Techniques and Applications*, Springer.

[10] Evans, W., Kenyon, C., Peres, Y. and Schulman L. J. (2000) Broadcasting on trees and the Ising Model, *Ann. Appl. Prob.*, **10 no. 2**, 410–433.

[11] Fitch, W. M. (1971). Toward defining the course of evolution: minimum change for a specific tree topology. *Syst. Zool.* **20**, 406–416.

[12] Hajek, B. and Weller, T. (1991). On the maximum tolerable noise for reliable computation by formulas. *IEEE Trans. on Information Theory* **37**(2), 388–391.

[13] Higuchi, Y. (1977). Remarks on the limiting Gibbs state on a (d+1)-tree. *Publ. RIMS Kyoto Univ.* **13**, 335–348.

[14] Ioffe, D. (1996a). A note on the extremality of the disordered state for the Ising model on the Bethe lattice. *Lett. Math. Phys.* **37**, 137–143.

[15] Ioffe, D. (1996b). A note on the extremality of the disordered state for the Ising model on the Bethe lattice. In *Trees*, B. Chauvin, S. Cohen, A. Roualt (Editor).

[16] Kenyon, C., Mossel, E. and Peres, Y. (2001). Glauber dynamics on trees and hyperbolic graphs, *Preprint.*

[17] Kesten, H. and Stigum, B. P. (1966) Additional limit theorem for indecomposable multidimensional Galton-Watson processes, *Ann. Math. Statist.* **37**, 1463–1481.

[18] Lyons, R. (1990) Random walks and percolation on trees. *Ann. Probab.* **18**, 931–958.

[19] Lyons, R. and Pemantle R. (1992) Random walk in a random environment and first-passage percolation on trees. *Ann. Probab.* **20,1** 125–136.

[20] Lovàsz, L. and Winkler P. (1998) Mixing times, in *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* **41**, 85–133.

[21] Mossel, E. (1998) Recursive reconstruction on periodic trees, *Random Structures and algorithms* **13,1** 81–97

[22] Mossel, E. (2001) Reconstruction on trees: Beating the second eigenvalue, *Ann. Appl. Probab.*, to appear.

[23] Pemantle, R. and Steif, J. E. (1999). Robust phase transitions for Heisenberg and other models on general trees. *Ann. Probab.* **27 no. 2**, 876–912

[24] Shamir, A. (1979). How to share a secret? Communications of the ACM **22** , 612–613.

[25] Spitzer, F. (1975). Markov random fields on an infinite tree. *Ann. Probab.* **3**, 387–394.

[26] Steel, M. (1989). Distributions in bicolored evolutionary trees. *Ph.D. Thesis*, Massey University, Palmerston North, New Zealand.

Elchanan Mossel
Microsoft research

1 Microsoft way
Redmond WA 98584
`mossel@microsoft.com`

Yuval Peres
Department of Statistics
367 Evans Hall, University of California
Berkeley, CA 94720-3860
`peres@stat.berkeley.edu`